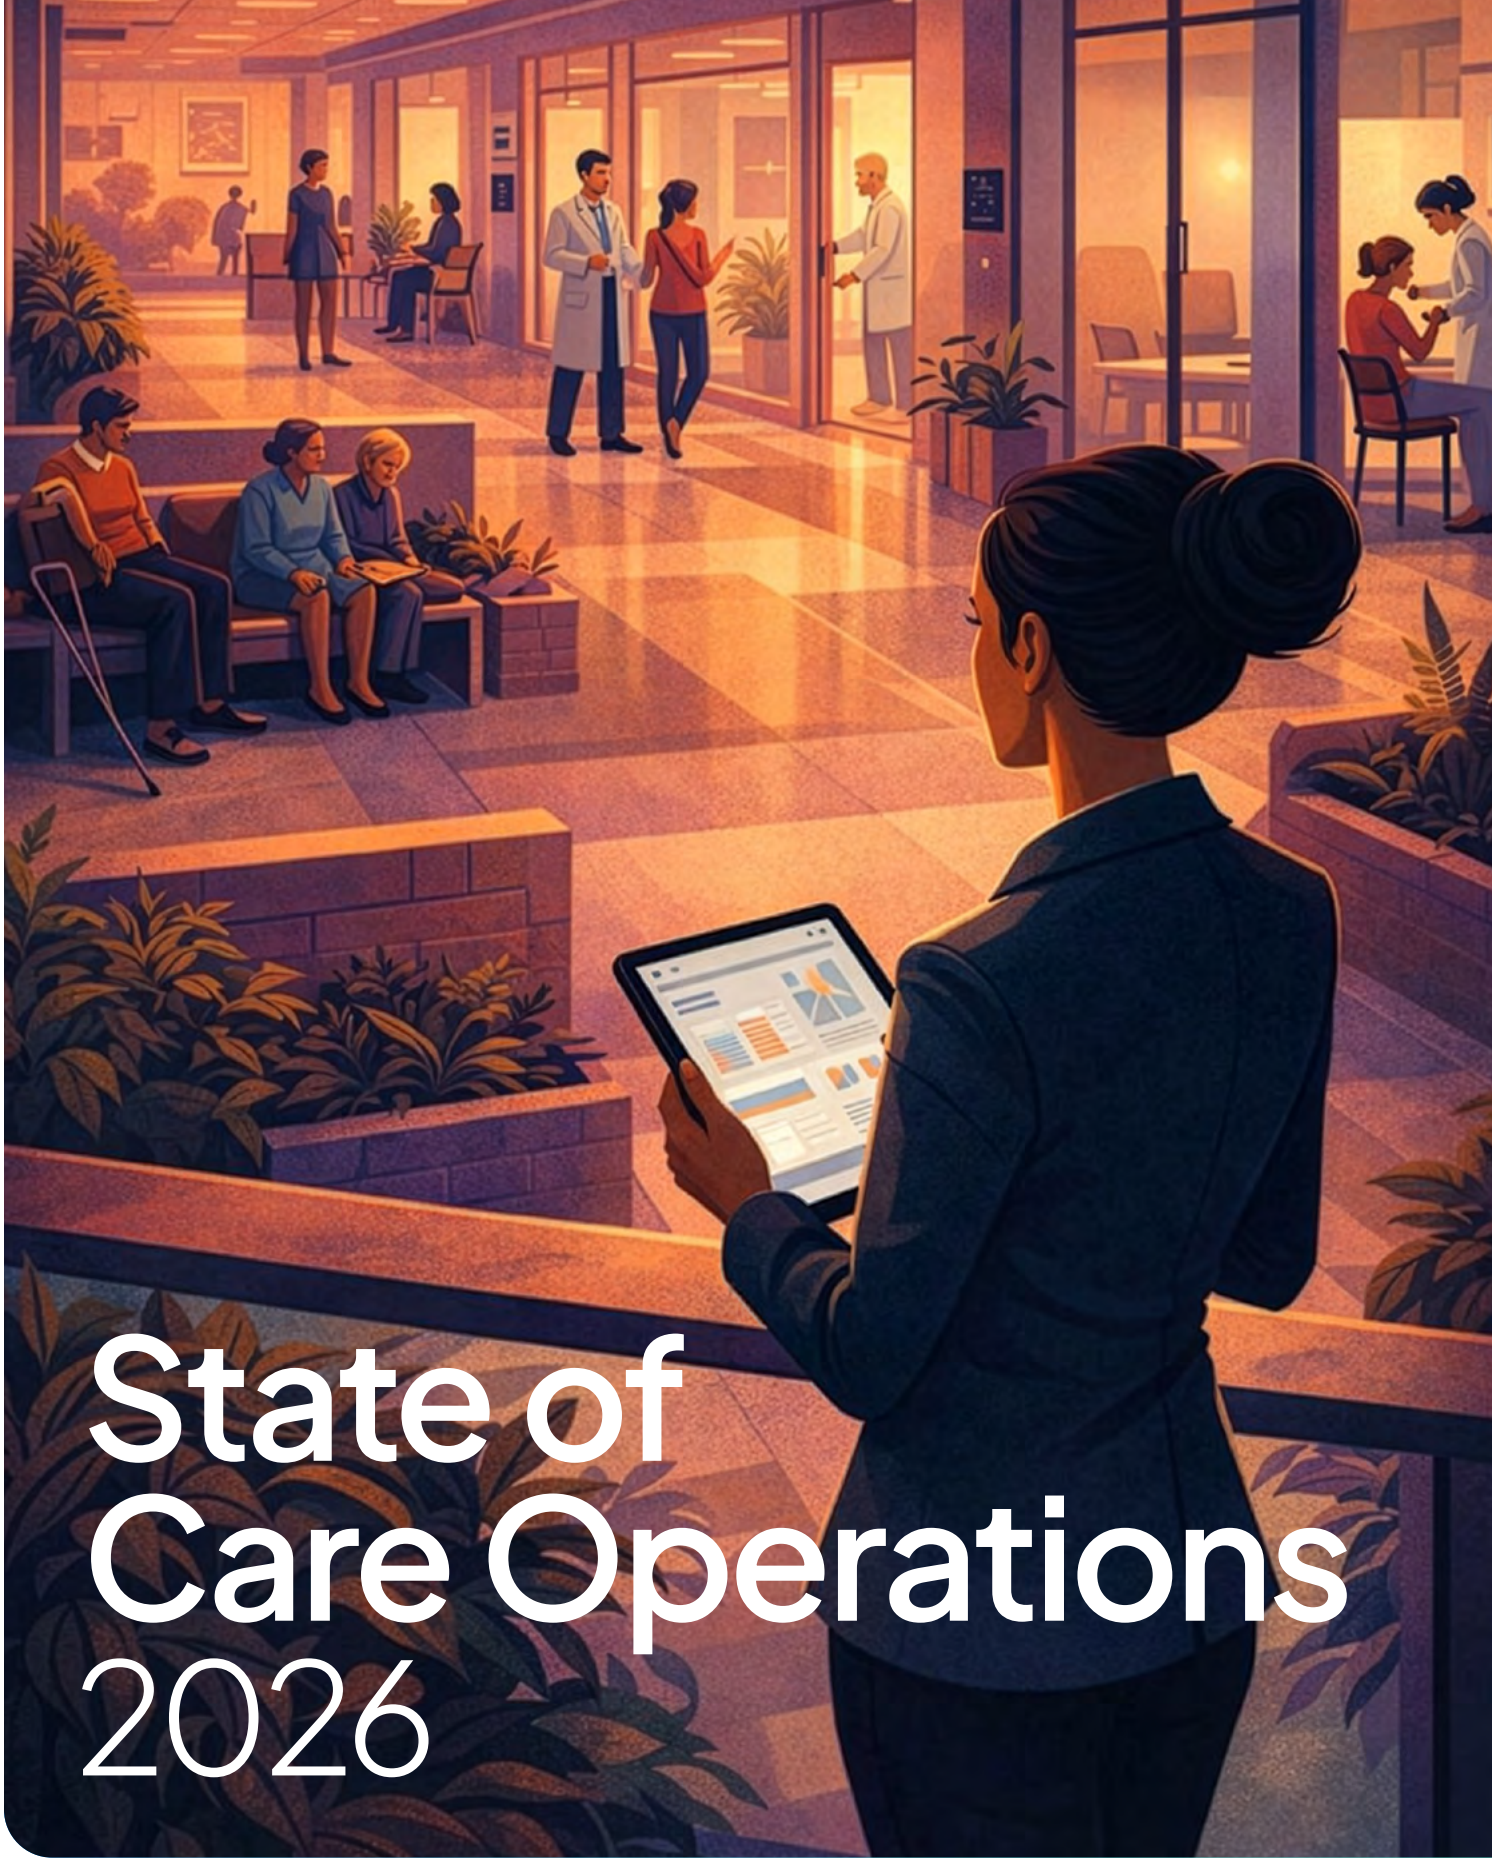

State of Care Operations 2026

EVIDENCE-BASED OPERATIONS

Stat

Introduction

Healthcare operations are under more pressure than ever—but clarity remains in short supply.

Across primary care, health centers, private practices, and hospitals, leaders are being asked to do more with fewer resources: fewer staff, tighter margins, more complex patients, and rising expectations around access and experience. In response, organizations have invested heavily in digital tools, satisfaction surveys, and isolated workflow optimizations. Yet despite these efforts, the lived experience of patients—and the daily frustration of frontline teams—has changed far less than expected.

> ***The problem is not a lack of effort. It is a lack of fundamental truth.***

Most industry benchmarks rely on lagging indicators (satisfaction surveys, throughput averages, utilization assumptions) or data sources shaped by human input and interpretation. These measures are valuable, but incomplete. *They tell us how people feel about care—not what actually happens inside the visit.*

This report closes that gap.

With this report, Stat provides an inside-the-visit view of care operations, built from passive, minute-level flow data captured directly from clinical environments. Rather than asking patients or staff to recall experiences after the fact, this analysis measures what truly occurs: when patients arrive, how long they wait, when rooms sit idle, and where time is lost.

The result is a clear, sometimes uncomfortable, picture of modern care delivery. Across more than one million patient visits, a consistent pattern emerges:

- Total visit length is driven far more by unmanaged flow than by care delivery
- Waiting room time and alone time dominate the patient experience
- Provider time is remarkably consistent and rarely the root cause of delay
- Capacity constraints are operational, not physical

The thesis of this report is simple:

Healthcare does not have a speed problem, a staffing problem, or a space problem. It has a flow problem.

Methodology & Scope

Data Collection

The findings in this report are based on passive, high-fidelity flow data captured without manual data entry.

Stat records minute-level location and timing signals throughout the patient visit, from arrival through departure. Because the data is collected automatically, it avoids the limitations common to self-reported metrics:

- No human bias
- No retrospective estimation
- No selective sampling
- No leadership reinterpretation

This approach enables consistent, auditable definitions of key ops moments—waiting, rooming, and provider interaction—across organizations and care settings.

Why This Matters

Without consistent definitions, improvement efforts drift toward anecdotes and gut feel. Two clinics may both report a "45-minute visit," yet deliver dramatically different patient experiences depending on where that time is spent.

This dataset allows true apples-to-apples comparison by measuring:

- Door-to-door cycle time
- Waiting room time
- Alone time in exam rooms
- Time with providers
- Time with staff
- Room utilization

Rather than focusing on perceptions of efficiency, the analysis focuses on how time is actually used.

Scope of Stat Data and Analysis

Who: Health centers, private practices, hospitals
What: Family Medicine, Internal Medicine, Pediatrics
Where: Alaska to Puerto Rico; New York to California
When: January 1 – December 31, 2025

Patient visits: 1,000,000+
Organizations: 36
Departments: 178

Flow & Cycle Time

Why Time is the Most Honest Measure of Care Operations

Cycle time is often discussed, rarely understood, and almost never measured well.

In many organizations, visit length is treated as a blunt performance metric—something to be shortened, justified, or explained away. But total cycle time, on its own, tells us very little. What matters is where that time is spent, when it accumulates, and why it persists.

This section examines care operations through the lens of time—as a sequence of lived moments inside the visit. By breaking the visit into its component parts, it becomes possible to distinguish value from waste, intentional buffering from accidental delay, and systemic design from individual behavior.

Flow and cycle time are not about speed. They are about coordination.

When flow is unmanaged, time expands invisibly. When flow is designed deliberately, capacity, experience, and sustainability improve together.

Total Cycle Time

The average primary care visit lasts just over 50 minutes from door to door. However, only a fraction of that time is spent delivering care.

Across the dataset, **approximately 62% of total visit time is consumed by waiting and alone time**, while only about one-third represents value-adding interaction with clinical staff.

From the patient's perspective, the visit is not defined by the 12 minutes spent with a provider—it is defined by the 30+ minutes spent waiting, alone, or uncertain about what comes next.

The system is optimized to protect provider schedules, not patient flow.

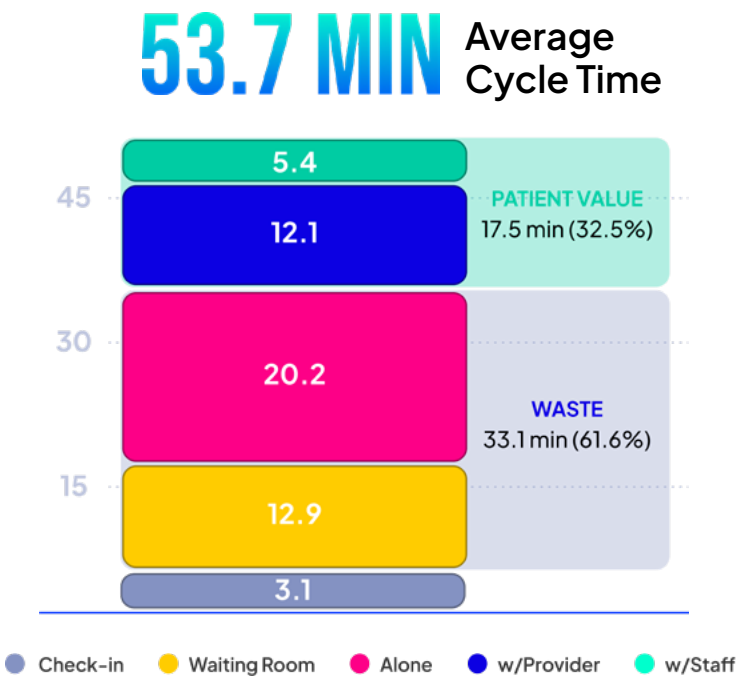

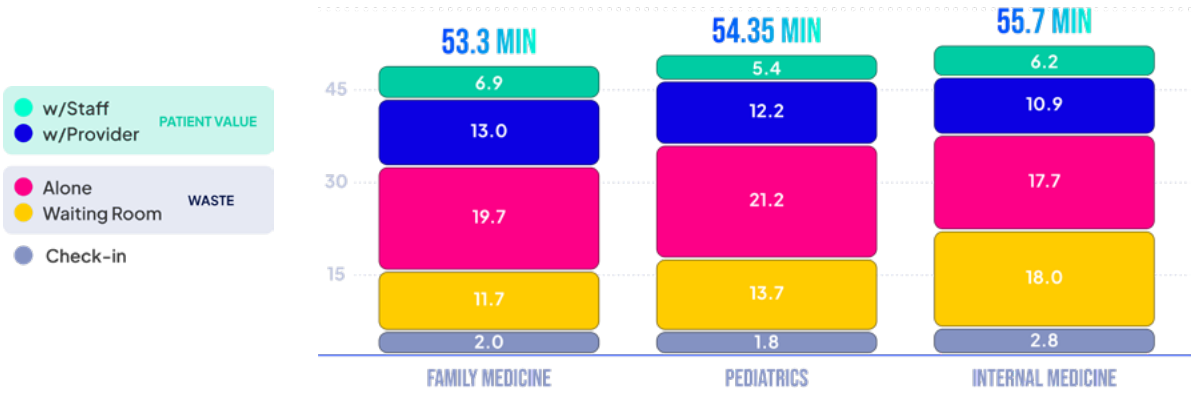

Cycle Time by Specialty

Examining cycle time across specialties—**Family Medicine, Pediatrics, and Internal Medicine**—reveals total visit lengths are nearly identical.

Despite differences in patient age, visit complexity, and staffing models, cycle times converge tightly in the mid-50-minute range. This consistency is not coincidence. It is evidence that the primary drivers of inefficiency live outside clinical care.

Workflow design dominates specialty differences. This means improvement strategies should be horizontal, not siloed. Optimizing rooming, handoffs, and discharge processes delivers more impact than specialty-specific tweaks.

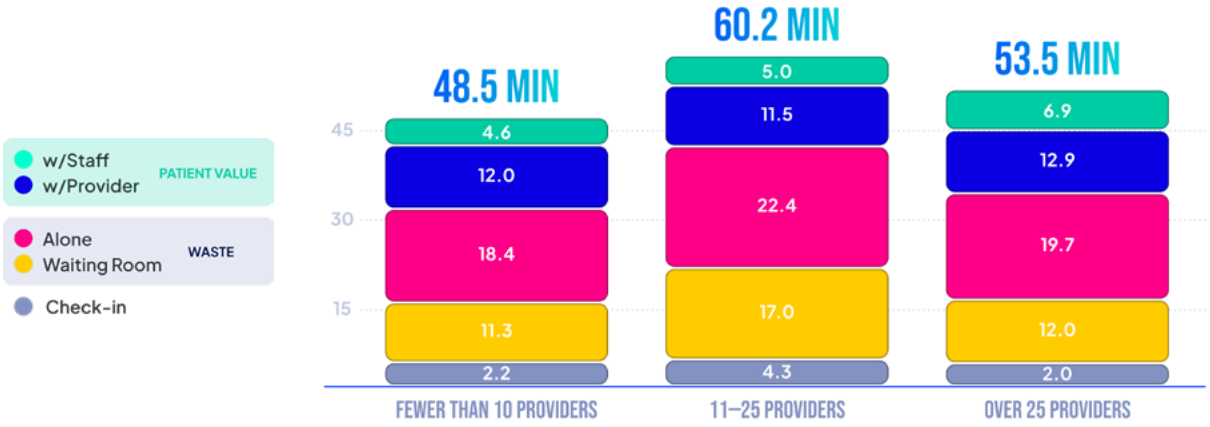

Cycle Time by Organization Size

Organization size introduces another revealing pattern.

- **Small practices** (<10 providers) exhibit the shortest cycle times
- **Mid-sized organizations** (11–25 providers) show the longest visits
- **Large organizations** (>25 providers) recover some efficiency through scale and standardization

Growth introduces complexity faster than coordination As organizations expand, informal communication breaks down before formal process matures. Alone time grows fastest in this middle stage, creating patient frustration and wasted capacity.

Size alone does not determine performance. Operational maturity does.

Check-In Time

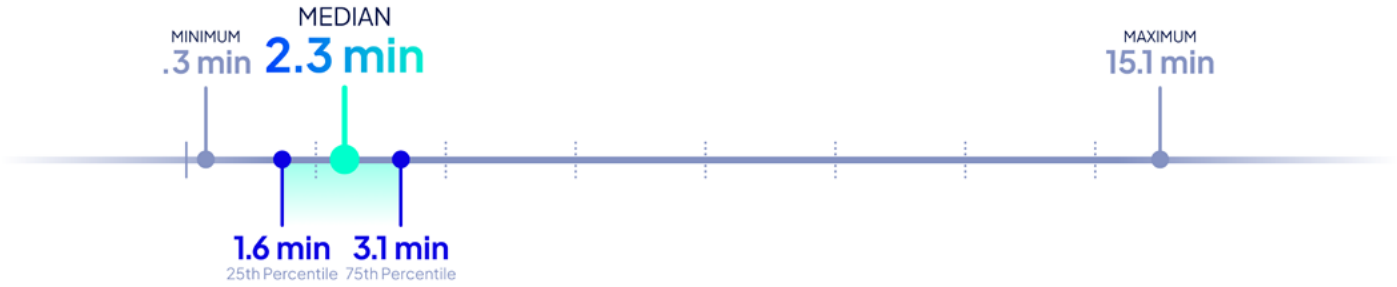

Front-desk workflows have largely been optimized.

Median check-in time across the dataset is just over two minutes, with a tight distribution for most visits. While outliers exist—often driven by insurance issues or first-time patients—registration is not where systemic friction lives.

This matters because it challenges a common improvement reflex: *fixing what is most visible rather than what is most impactful.*

Waiting Room

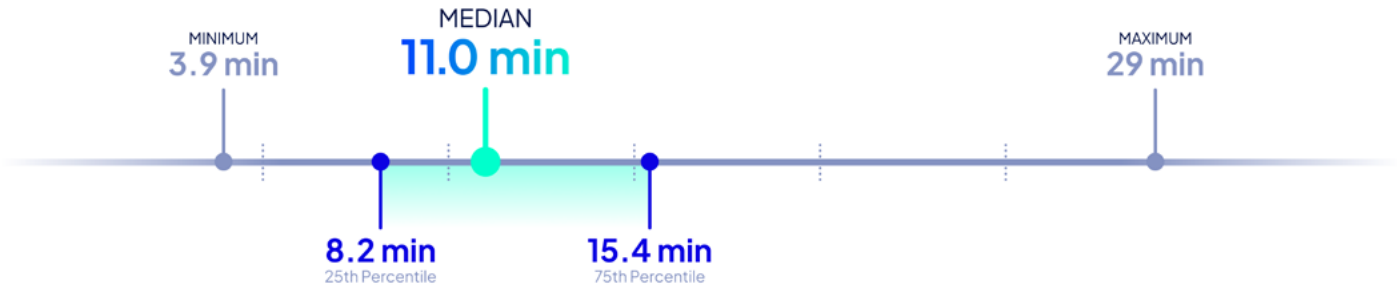

Waiting room time tells a different story. While median waits are moderate, the upper quartile increases sharply, with maximum waits approaching 30 minutes. These delays are not by accident, nor are they random. *They reflect mismatches between scheduling, staffing, and room availability.*

The waiting room absorbs operational variance. Because leaders can see patients waiting, attention often focuses here—even though this is only part of the problem.

Alone Time

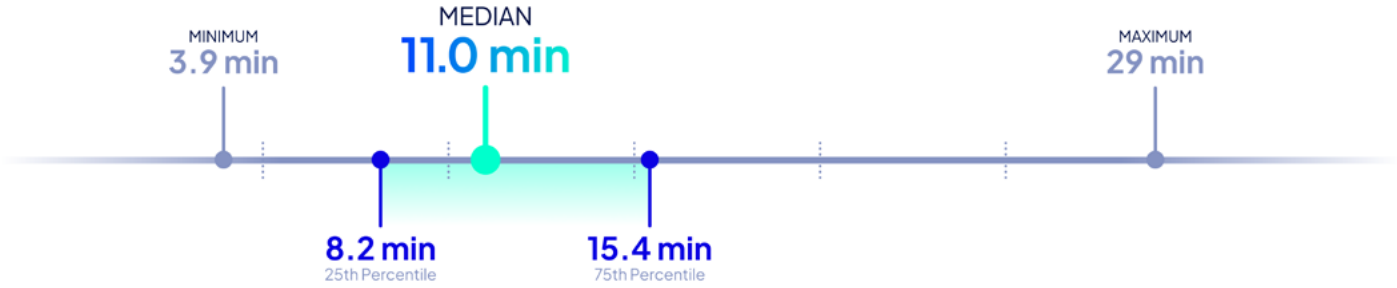

The single largest block of wasted time occurs after the patient is roomed.

Median alone time in exam rooms exceeds 18 minutes, with the upper quartile stretching beyond 24 minutes and extreme cases surpassing 30 minutes.

Early rooming creates the illusion of efficiency while generating significant hidden waste. Patients are ready long before the system is. This buffering strategy protects provider schedules, but it does so at the expense of patient experience and room utilization.

Time with Provider

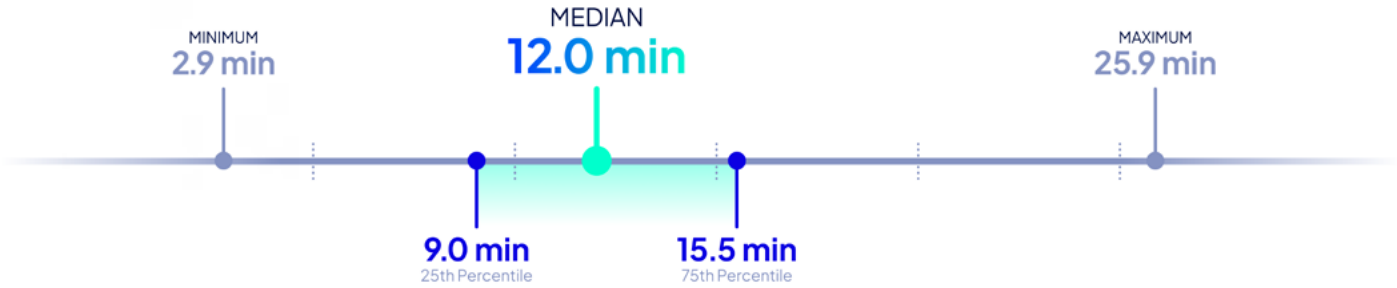

Time spent with providers is remarkably consistent across organizations.

Median provider interaction hovers around 12 minutes, with limited variation even at the upper end. This directly contradicts the belief that long visits are caused by providers "taking too long."

Providers operate within predictable time bands. The delays happen between them.

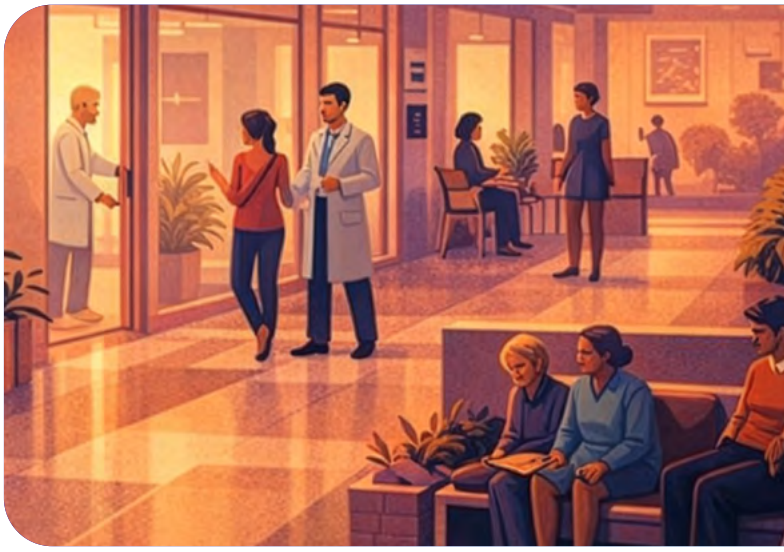

Time with Staff

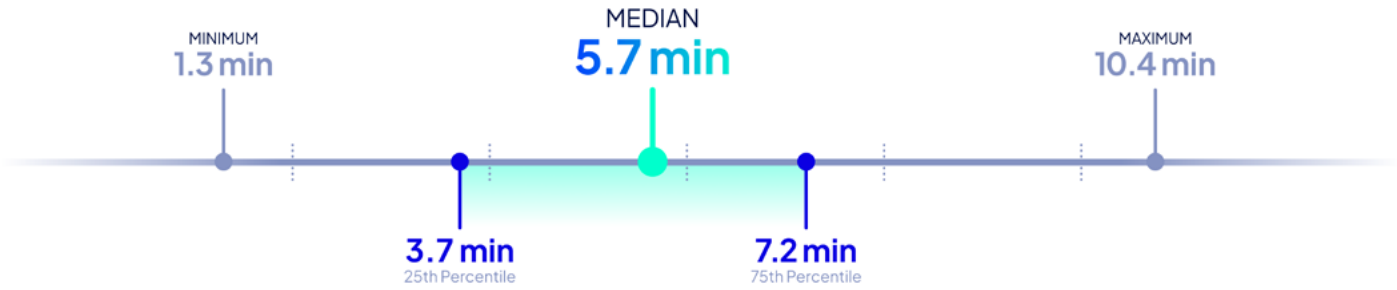

Time spent with clinical staff tells a simple, yet complementary, story.

Median staff interaction time is relatively low, with a narrow interquartile range. When staff time expands, it is typically driven by interruptions, rework, or waiting for downstream steps—not slow execution.

Staff are not inefficient. They are often disconnected from real-time flow.

Myth #1: The "Lobby" Problem

> *"Our patients are frustrated because they wait too long in the lobby."*

When waiting room time, provider time, and alone time are compared directly, a critical insight emerges: **patients lose more time alone in exam rooms than anywhere else.**

The industry tends to over-index on lobby wait times because they are visible. Exam room delays remain hidden behind closed doors.

Healthcare does not have a lobby problem. It has an invisible waiting problem.

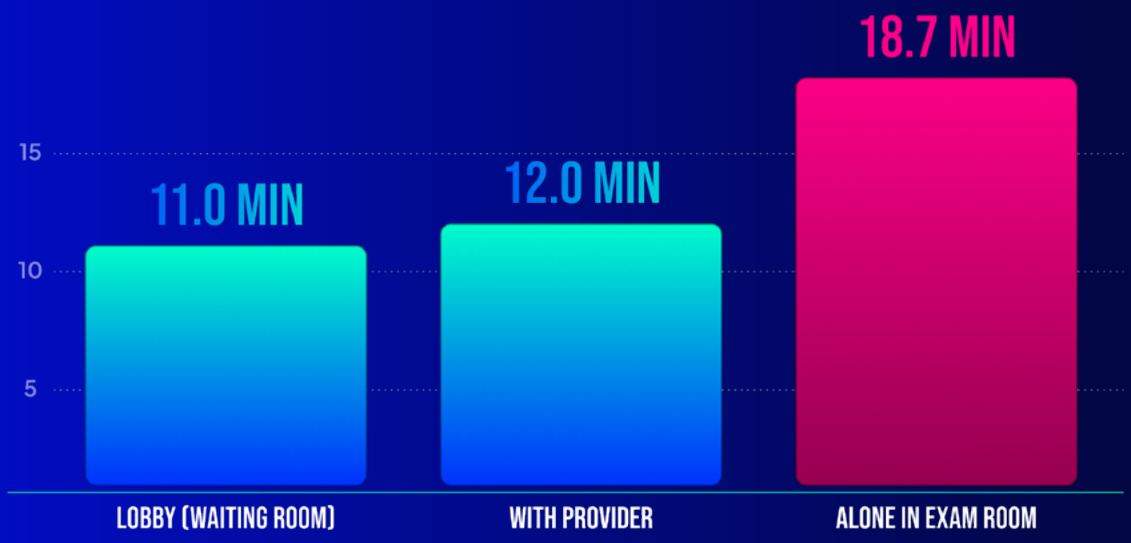

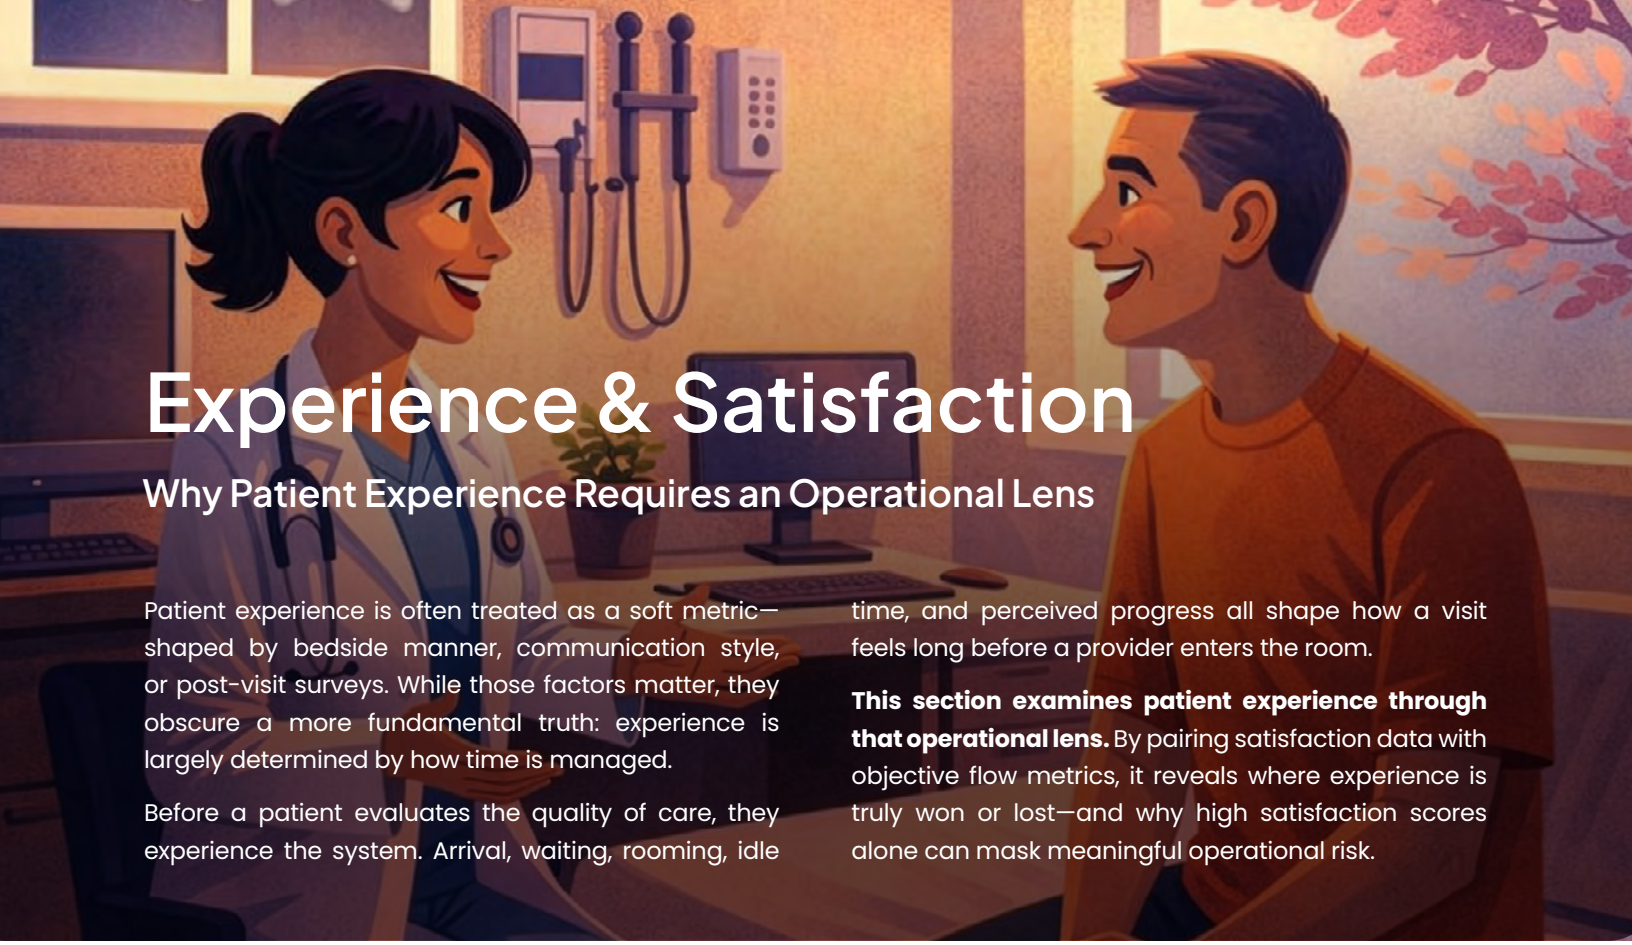

Experience & Satisfaction

Why Patient Experience Requires an Operational Lens

Patient experience is often treated as a soft metric—shaped by bedside manner, communication style, or post-visit surveys. While those factors matter, they obscure a more fundamental truth: experience is largely determined by how time is managed.

Before a patient evaluates the quality of care, they experience the system. Arrival, waiting, rooming, idle time, and perceived progress all shape how a visit feels long before a provider enters the room.

This section examines patient experience through that operational lens. By pairing satisfaction data with objective flow metrics, it reveals where experience is truly won or lost—and why high satisfaction scores alone can mask meaningful operational risk.

Patient On-Time %

Patient arrival patterns are remarkably consistent.

Across the dataset, patient arrival times form a near-normal distribution centered around the scheduled appointment time. Most patients arrive within a narrow window of their appointment—neither excessively early nor meaningfully late.

This matters because it challenges a persistent operational narrative: that patient punctuality is a primary driver of delays.

The data shows otherwise.

Normal arrival variation is not a failure of patients; it is an expected condition that operational systems must absorb. When schedules are brittle—tightly packed without buffer or flow awareness—even small variations cascade into waiting, idle rooms, and missed capacity.

> ***Patients are not late.***
> ***Schedules are fragile.***

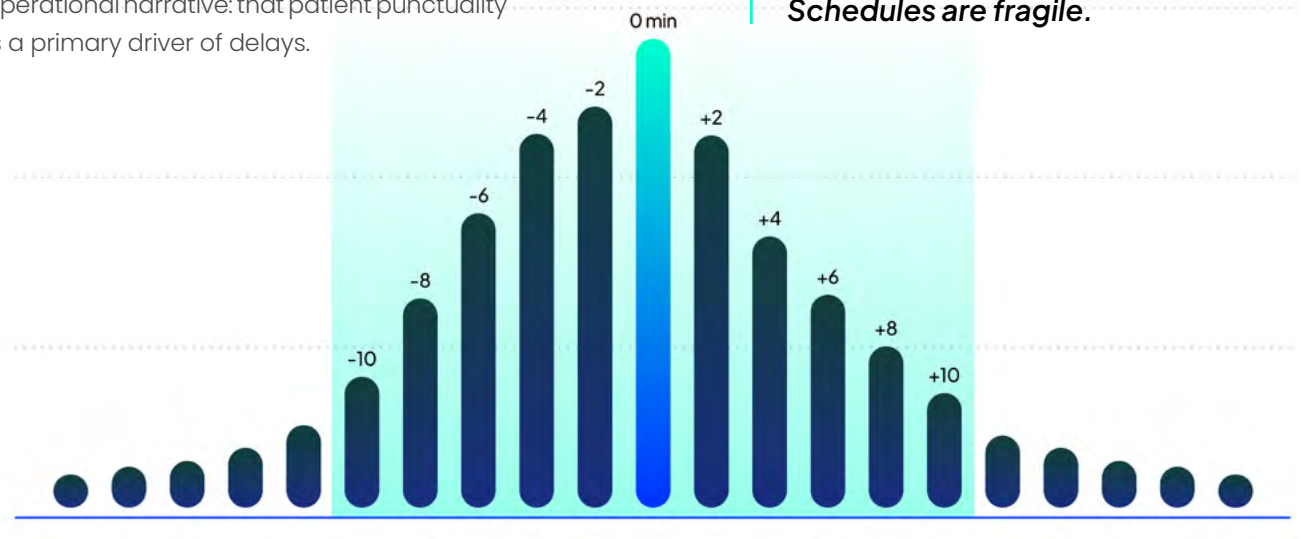

Top-Box Patient Satisfaction

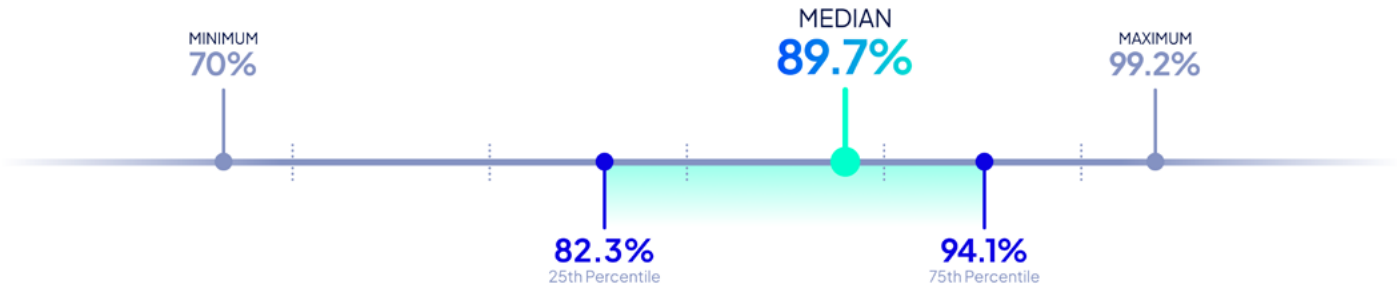

Overall patient satisfaction across visits is high. The median top-box satisfaction score approaches 90%, with a relatively narrow distribution across organizations. Even the lower end of performance remains well above what many industries would consider acceptable.

At first glance, this appears reassuring. But when viewed alongside flow and cycle time data, it reveals something more concerning.

High satisfaction can coexist with significant operational waste. *Patients adjust expectations.* Over time, long waits and inefficiencies become normalized—especially when clinical interactions remain positive. Satisfaction scores reflect this adaptation, not necessarily operational excellence.

What Drives Patient Satisfaction?

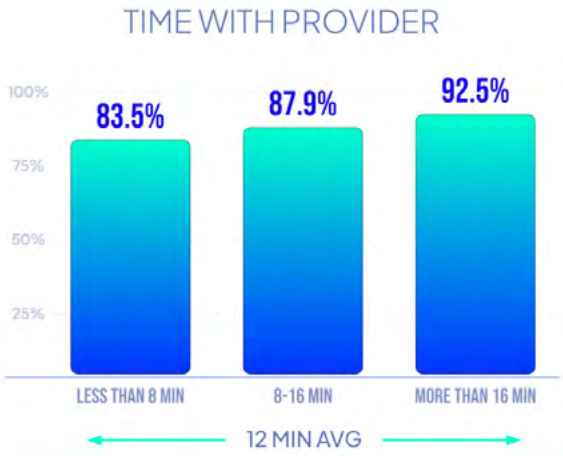

When satisfaction is segmented by operational factors, two drivers stand out clearly:

1. Waiting time
2. Time with the provider

Patients who experience shorter waits report meaningfully higher satisfaction. Likewise, adequate provider interaction strongly correlates with positive perception of the visit.

Extremely long waits disproportionately harm satisfaction, while incremental improvements below already-acceptable thresholds deliver diminishing returns. In other words, eliminating the worst delays matters far more than optimizing already-good averages.

> ***One bad wait outweighs five good visits in patient memory.***

Utilization & Throughput

Why Capacity Feels Scarce—Even When It Isn't

Capacity is one of the most emotionally charged topics in care operations.

When clinics feel busy, the default assumption is that demand has outpaced resources. Providers feel rushed, staff feel stretched, and leaders conclude that the organization has hit a physical limit—too few rooms, too few people, too many patients.

But perceived capacity and actual capacity are not the same thing.

This section examines utilization and throughput through an operational lens, using objective data to separate felt congestion from true constraint. By measuring how often rooms are actually used, when they sit idle, and how throughput varies across organizations, it reveals why many clinics feel maxed out long before they are—and where meaningful capacity is hiding in plain sight.

Exam Room Utilization

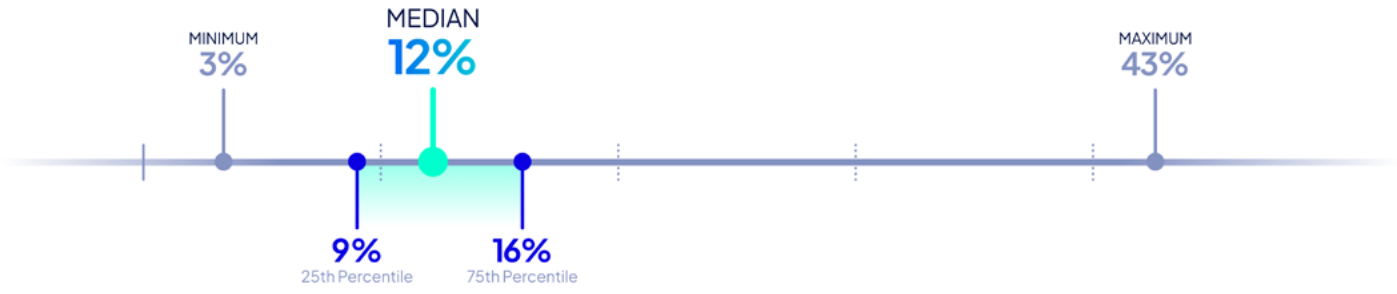

Clinics often feel full.

Schedules are packed, staff are stretched, and providers feel constant pressure to move faster. The natural conclusion is that capacity has been reached—that more rooms, more staff, or longer hours are required to keep up with demand.

The data tells a very different story.

Across the dataset, median exam room utilization is just 12%, with a 25th percentile near 9% and a 75th percentile of 16%. Even the highest-performing sites rarely exceed 43% utilization.

In other words, the average exam room sits empty nearly nine out of every ten minutes it is available.

This is not a space problem. It is a flow problem.

Room Use by Day

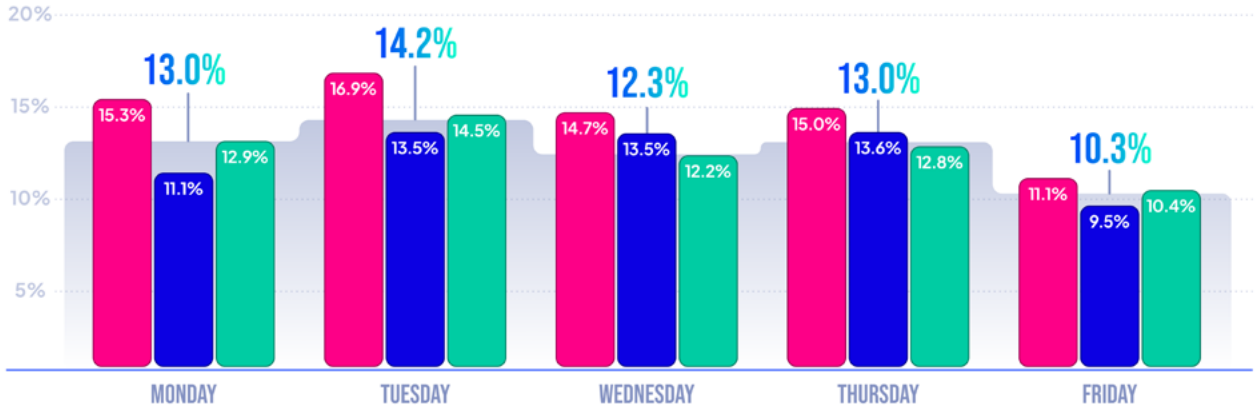

Exam room utilization is not evenly distributed across the week. Midweek days, particularly Tuesdays, show the highest room utilization across organizations. In contrast, Mondays and Fridays tend to operate well below peak capacity.

This variation is not driven by patient demand alone. It reflects scheduling conventions, provider availability patterns, and long-standing operational habits.

The implication is critical: perceived capacity constraints are often the result of when visits are scheduled, not how many visits an organization can support. By smoothing volume across the week—rather than concentrating visits into preferred days—organizations can unlock meaningful capacity without adding rooms, staff, or hours.

Room Use by Hour

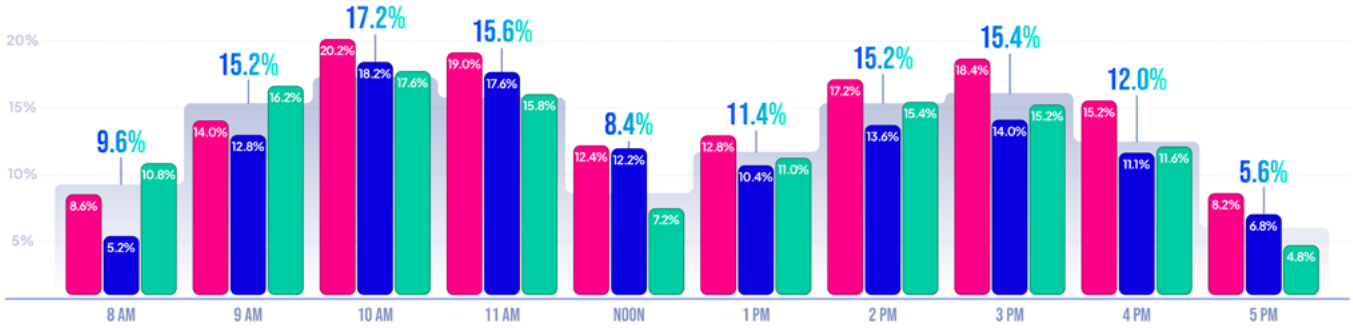

When room utilization is examined by hour of day, clear patterns emerge.

- Mornings consistently show the highest room utilization
- Midday peaks reflect batching and congestion, not sustained demand
- Early mornings and late afternoons remain underutilized across organizations

These patterns are remarkably consistent across sites.

They reveal that utilization follows scheduling habits—not patient need. Capacity exists, but it is unevenly accessed.

> ***When visits are stacked into the same windows, rooms feel scarce even when they are mostly empty.***

Myth #2: The "Space" Problem

> *"We are out of space and need to build more exam rooms."*

Because underutilization is invisible in real time, leaders often interpret congestion as a shortage of space. *The data directly challenges this assumption.*

Even on the busiest days, exam room utilization rarely approaches levels that would justify physical expansion. Building more rooms would simply increase the number of idle assets—not improve throughput.

You don't fix flow by adding space.

13.9%

AVERAGE EXAM ROOM UTILIZATION

Visits per Provider Hour

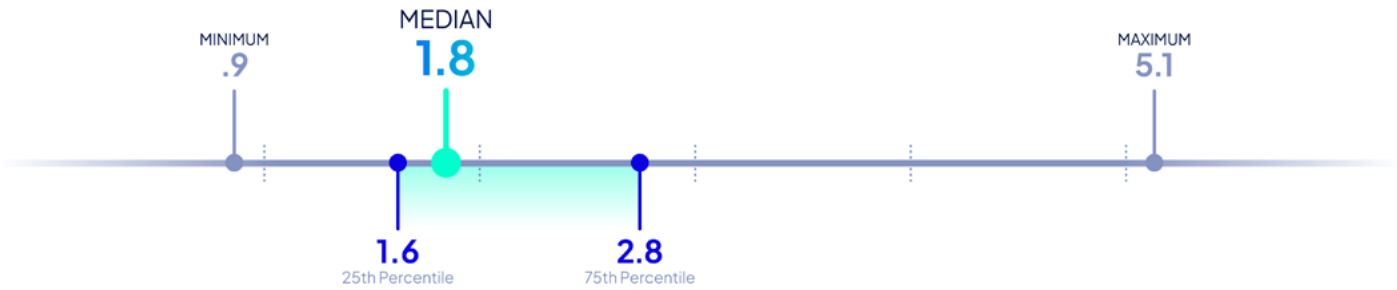

Visits per provider hour vary dramatically across organizations. High performers do not move faster or shorten visits. Instead, they:

- Reduce idle gaps between patients
- Ensure the next patient is ready when the provider is available
- Minimize interruptions and rework

Throughput emerges as an operational outcome, not a behavioral expectation.

Providers don't create productivity. Systems enable it.

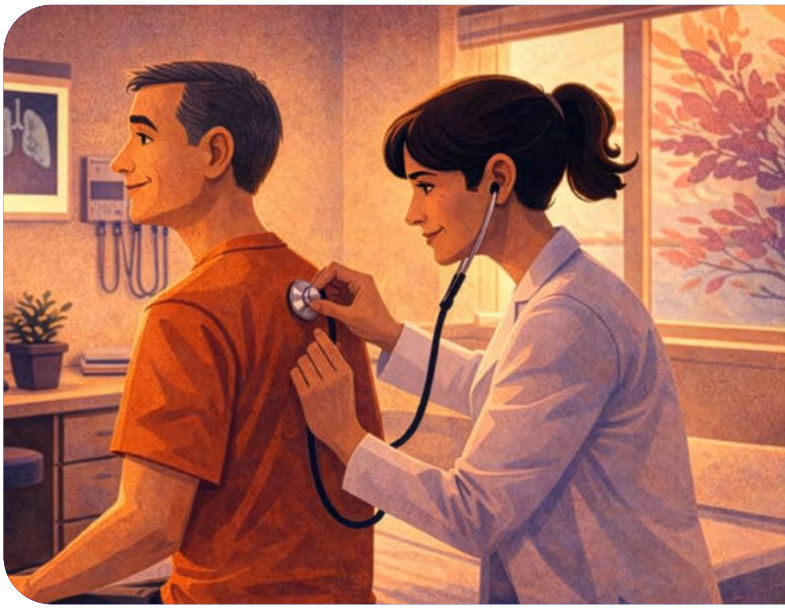

Cycle Time and Productivity

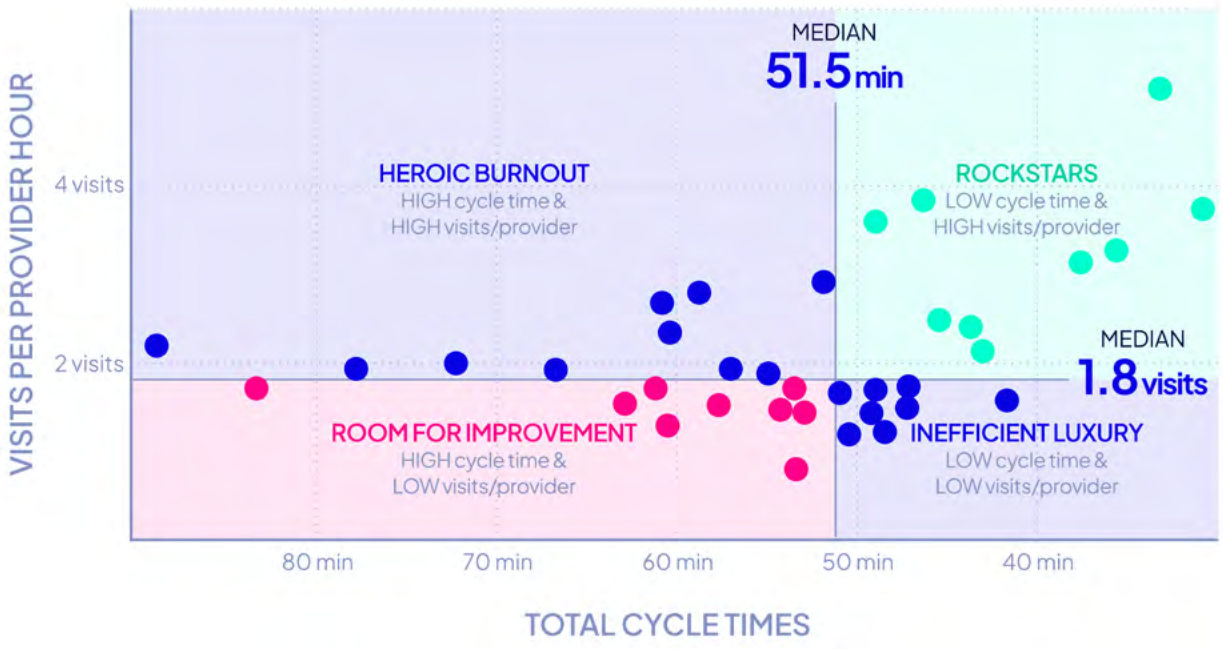

When cycle time and throughput are viewed together, a critical insight emerges: low cycle time does not require low productivity. Inefficient organizations cluster into three distinct patterns:

1. **Room for Improvement:** inefficiency experienced as long visits desite having low visits per provider
2. **Heroic Burnout:** high throughput achieved by providers powering through long cycle times
3. **Inefficient Luxury:** inefficiency through low throughput, despite short cycle times

The highest performers operate in the **Rockstars** quadrant—delivering more visits per provider hour while maintaining shorter visits.

Rockstars do not rush care. They eliminate waste.

Care Operations Implications for 2026

The data in this report points to a clear conclusion: the next gains in care operations will not come from working harder, moving faster, or building more.

They will come from managing flow.

As organizations plan for 2026, four operational priorities stand out.

Treat Flow as a First-Order Need

Flow cannot be owned indirectly or managed part-time. High-performing organizations assign explicit accountability for the end-to-end visit—not just individual steps like check-in, rooming, or scheduling. When someone is responsible for the whole experience, tradeoffs become visible and delays become solvable.

If flow is everyone's job, it is no one's job.

Design for Variability, Not Perfection

Patients arrive with normal variation. Visits take different amounts of time. Staff availability changes throughout the day. Systems that assume perfect execution break under real-world conditions.

Operational resilience comes from designing schedules, rooming practices, and staffing models that absorb variability rather than amplify it. Eliminating extreme delays delivers far more value than optimizing already-acceptable averages.

Unlock Existing Capacity Before Adding More

Most organizations are operating far below their true capacity. Before adding rooms, staff, or hours, leaders should ask a simpler question: *Are we using what we already have?*

Improving flow unlocks idle rooms, reduces provider gaps, and increases throughput without increasing burnout or cost. Expansion without utilization discipline simply scales inefficiency.

Measure What Actually Happens Within Your Organization

You cannot manage what you cannot see. Lagging indicators like satisfaction and productivity summaries are useful, but insufficient. Organizations that improve fastest rely on objective, inside-the-visit data to guide decisions and validate change.

Measurement is not about surveillance. It is about operational clarity.

What Gets Measured, Gets Improved

The insights in this report are only possible because they are grounded in objective, inside-the-visit data—the kind of data most organizations have never been able to see, let alone act on.

By passively capturing minute-level flow data without human input, Stat gives care operations leaders a clear, auditable view of how time and space are affected as people move through their systems. This clarity is what enables evidence-based operations: decisions rooted in reality, improvement efforts focused on the true constraints, and changes that can be measured, validated, and sustained.

To understand how these insights apply inside your own clinics—and where your biggest opportunities are hiding—schedule a demo with Stat and see your operations as they truly are.

Schedule a Demo stat.io/demo sales@stat.io